

Remarks

Original Claims 1-43 are currently pending in this application.

Claim Objections

5 **Claims 10 and 18** are objected to on the ground that “an Mx record” should read
“a Mx record.” The Applicant appreciates the Examiner’s suggestion for this
modification. However, the Applicant believes that the claim language is correct and
wishes to avoid making any unnecessary claim amendments that might be improperly
interpreted as being required for allowance. As written, Claims 10 and 18 meet the
10 requirements of §112.

In the current office action the Examiner states “[a]ppropriate correction is
required.” The Applicant requests that if the Examiner is going to require the above
correction, that the Examiner provide a statutory basis for doing so.

The Applicant notes that the entry for “mx record” within www.wikipedia.org
15 starts “An MX record or Mail exchange record.”

Rejections under 35 USC §103

**Claims 1-9 and 11-42 are rejected under 35 U.S.C. 103 as being unpatentable
over Fletcher et al. (USPN: 2002/0178238), hereinafter Fletcher, in view of Singh et
al. (USPN: 2004/0013113), hereinafter Singh.**

20 **Regarding Claim 1,**

Claim 1 recites:

*1. A caching server comprising:
an answer cache configured to access answer information through a flat data
structure;
25 a referral cache configured to store referral information; and*

computer instructions configured to translate a domain name into DNS information by examining the answer cache and, responsive to the results of examining the answer cache, examining the referral cache.

5 In rejecting Claim 1, the Examiner states:

10 “Fletcher teaches a caching server comprising an answer cache (i.e. the cache which stores the address information that is local within the terminal) configured to store answer information ... a referral cache (i.e. the cache which stores referral information to forward the query to the remote terminal’s answer cache across the communication network) configured to store referral information (i.e. the referral information) ... (e.g. see paragraph [0008]).”

The Applicant traverses this statement.

15 First, the caching servers of Fletcher do not appear to teach a caching server including both “*an answer cache*” and a separate “*referral cache*.” Fletcher does teach a cache 405 as illustrated in FIG. 4A of Fletcher. However, this is one cache and, thus, does not teach both “*an answer cache*” and a separate “*referral cache*.” The Examiner’s statement refers to “the cache which stores referral information.” However, the Applicant is unable to find any such teaching within the cited art. Fletcher appears to be
20 silent as to where referral information is stored. As such, referrals may not be stored in a cache at all, or may be stored in the same single cache as the answer information. In neither of these two cases would the limitations of Claim 1 be taught. The Applicant, therefore, requests that the Examiner specifically point out teaching of two separate caches within a caching server, or allow Claim 1, and those claims that depend therefrom.

25 Second, assuming for the sake of argument that the cache illustrated in FIG. 4A of Fletcher teaches an answer cache, Fletcher does not appear to teach, “*a referral cache configured to store referral information*,” as recited in Claim 1. Fletcher does teach:

[i]n the iterative resolution, a name server that receives the query simply gives the best answer it already knows back to the querying name server ... The name

server that is queried consults its local database and/or cache for the requested data. If the queried name server does not find the data, the queried name server makes its best attempt to provide the querying name server with data that helps the querying name server in continuing the resolution process by issuing
5 “referrals,” (paragraph [0030]).

However, while this text does teach “referrals” it does not appear to teach that the referrals are stored in a “*referral cache*,” as distinguished from a separate “*answer cache*.” This text of Fletcher is silent as to where referrals are stored or how they are
10 obtained. As such, Fletcher merely teaches that referrals are issued as a “best answer” by a name server if other data is not found (paragraph [0030] above).

Fletcher also teaches a cache that is characterized as “store[ing] symbolic address and numeric addresses to provide mapping from a symbolic address to a numeric address and vice-versa,” (paragraph [0046]). Assuming for the sake of argument that this is a
15 teaching of answer information stored in the cache of FIG. 4A, this teaching would characterize the cache 405 of Fletcher as an “*answer cache*” rather than a “*referral cache*.” Thus, the cache 405 of Fletcher would not be understood as teaching a “*referral cache*.”

In contrast, Claim 1 recites both “*an answer cache*” and “*a referral cache*
20 *configured to store referral information*.” These are two different caches characterized by the specification as filed as being separate (Title) and the “*referral cache*” is configured to store a specific type of information. The Applicant is unable to find “*a referral cache configured to store referral information*,” in the cited art. The Applicant, therefore, requests that the Examiner specifically point out teachings of “*a referral cache*
25 *configured to store referral information*,” within the cited art or allow Claim 1.

The Examiner appears to be inferring a separate referral cache from the teaching of “referrals.” It is the position of the Applicant that the separate “*referral cache*” and “answer cache” recited in Claim 1 cannot be inferred from the teaching of “referrals.” Such an inference would be improper because the referrals and other “best answers” may
5 be stored in the same cache. This would eliminate the need to search and manage two different caches.

MPEP §2112 provides that “[t]he fact that a certain result or characteristic *may* occur or be present in the prior art is not sufficient to establish the inherency of that result or characteristic” citing *In re Rijckaert*, 9 F.3d 1531, 1534, 28 USPQ2d 1955, 1957 (Fed.
10 Cir. 1993) (emphasis added). Further, “In relying upon the theory of inherency, the examiner must provide a basis in fact and/or technical reasoning to reasonably support the determination that the allegedly inherent characteristic necessarily flows from the teachings of the applied prior art.” *Ex parte Levy*, 17 USPQ2d 1461, 1464 (Bd. Pat. App. & Inter. 1990) (emphasis in original).

15 As stated above, a separate “*referral cache*” does not necessarily flow from a teaching of “referrals” because the referrals may be stored in a place other than a separate “*referral cache*,” e.g., an answer cache. It is, therefore, the position of the Applicant that the inference made by the Examiner is improper and that the cited art does not teach all of the limitations of Claim 1 even in combination. On at least this basis, the Applicant
20 requests that the Examiner allow Claim 1, those claims that depend therefrom, and Claims 11, 13, 17, 26, 34 and 38-39, which are rejected based on the same rationale as the rejection of Claim 1.

Second, Claim 1 includes the limitations “*examining the answer cache and, responsive to the results of examining the answer cache, examining the referral cache.*”

As pointed out above, it is the position of the Applicant that the cited art does not teach a separate “*referral cache.*” It is further the position of the Applicant that the cited art does
5 not teach a separate examination of “*a referral cache.*” The text cited by the Examiner teaches that “referrals” may be issued by a name server as the result of not finding “requested data” in a query (paragraph [0030]). Assuming for the sake of argument that the query of paragraph [0030] teaches, “*examining the answer cache,*” there does not appear to be any explicit teaching within Fletcher of performing a separate examination
10 of a referral cache. The Applicant, therefore, requests that the Examiner specifically point out a teaching of a separate examination of a referral cache, or allow Claim 1.

If the Examiner’s position is that a separate search is inherent to the teachings of paragraph [0030] of Fletcher, then the Applicant respectfully points out that referrals may be issued without a separate examination and a separate examination, therefore, does not
15 “necessarily flow” from the cited art. “In relying upon the theory of inherency, the examiner must provide a basis in fact and/or technical reasoning to reasonably support the determination that the allegedly inherent characteristic necessarily flows from the teachings of the applied prior art.” Ex parte Levy, *supra*. For example, a referral may be found during the initial query of paragraph [0030]. The query may return either “the
20 requested data” or the “referral” from a same cache, and thus not require a separate examination of a referral cache. Alternatively, a referral may be coded into software of the name server and, thus, be accessible without an examination of a referral cache.

Therefore, a separate examination is not inherent to the teachings of paragraph [0030] of Fletcher and the Examiner has not made a *prima facie* case for the rejection of Claim 1.

On at least these bases, the Applicant requests that the Examiner allow Claim 1, those claims that depend therefrom, and Claims 11, 13, 17, 26, 34 and 38-39, which are
5 rejected based on the same rationale as the rejection of Claim 1.

Third, in rejecting Claim 1, the Examiner admits, “Fletcher does not clarify that the answer cache stored answer information in a flat data structure,” and suggests that Singh teaches these limitations. Specifically, the Examiner states “Singh ... teaches that it is beneficial to use the hash table (i.e. the flat data structure) to perform the interface ID
10 (i.e. the answer information) lookup because it reduces the number of memory accesses and as a result of that, it is faster than the lookup in the tree structure (e.g. see paragraph [0050]).” The Applicant traverses this statement.

To establish a *prima facie* case of obviousness, three basic criteria must be met. First, there must be some suggestion or motivation, either in the references themselves or
15 in the knowledge generally available to one of ordinary skill in the art, to modify the reference or to combine reference teachings. Second, there must be a reasonable expectation of success. Finally, the prior art reference (or references when combined) must teach or suggest all the claim limitations. The teaching or suggestion to make the claimed combination and the reasonable expectation of success must both be found in the
20 prior art, and not based on applicant's disclosure. In re Vaeck, 947 F.2d 488, 20 USPQ2d 1438 (Fed. Cir. 1991). See MPEP § 2143 - § 2143.03.

As a motivation to combine Fletcher and Singh, the Examiner states “it would have been obvious to one ordinary skilled in the art at the time of the current invention

was made to implement the flat data structure in the answer cache to achieve the benefits described above.” The Applicant traverses this statement. The Examiner suggests that the combination would be obvious because an advantage would be achieved. The Applicant respectfully points out that the requirements of a *prima facie* case under §103

5 include “some suggestion or motivation” for the combination, not that the combination would produce an advantage. An invention can be non-obvious while still producing an advantage. The benefit must be apparent without the benefit of hindsight. It is, therefore, the position of the Applicant that the Examiner has not met the first element of a *prima facie* case for obviousness cited above. Specifically, the Examiner has not provided

10 “some suggestion or motivation, either in the references themselves or in the knowledge generally available to one of ordinary skill in the art, to modify the reference or to combine reference teachings” as required to establish a *prima facie* case under §103.

Further, in contrast with the Examiner’s suggested combination, Singh specifically teaches away from the inclusion of more than a mere interface ID in a hash

15 table. For example, in paragraph [0042] Singh teaches it “may be impractical to apply a ... hashing technique to the host portion of an IPv4 address.” This specifically teaches away from the storage of the DNS “address information” of Fletcher within a flat data structure. An answer to a query, in the context of Fletcher, that did not include host information would appear to be useless.

20 Further, it is the position of the Applicant that the teachings of Singh could not be combined with the teachings of Fletcher to produce a workable system, as suggested by the Examiner. Singh teaches that *part* of a destination address may be stored in a hash table. Specifically, in paragraph [0044], Singh teaches “[t]he destination address in this

example includes a TLA ID=101, a NLA ID, a SLA ID 112 and an interface ID=1743.”

In the paragraph cited by the Examiner ([0050]), Singh teaches that the interface ID=1743 can be stored in a hash table. Thus, the teachings of Singh are limited to a teaching that a part of a destination address, i.e., the interface ID, can be included in a
5 hash table.

A mere interface ID is not the equivalent of “*answer information*” as recited in Claim 1. As would be understood by one of ordinary skill in the art, “*answer information*” includes, for example, an IP address corresponding to a domain name, not merely an interface ID. An IP address is typically 128 bits while, according to Singh, an
10 interface ID is only 64 bits (FIG. 3). A cache that only returned a 64 bit value would be incompatible with standard IP protocols. Therefore, the combination suggested by the Examiner would result in an unworkable system. It is, therefore, the position of the Applicant that the Examiner has not met the second element of a *prima facie* case for obviousness, i.e., “a reasonable expectation of success.”

15 Finally, as pointed out above, the suggested combination does not suggest all the limitations of Claim 1. Therefore, it is the position of the Applicant that the Examiner has also not met the third requirement for a *prima facie* case of obviousness.

For at least these reasons, the Application requests that the Examiner allow Claim 1, those claims that depend therefrom, and Claims 11, 13, 17, 26, 34 and 38-39, which
20 are rejected based on the same rationale as the rejection of Claim 1.

Regarding Claims 11, 13, 17, 26, 34 and 38-39,

The Examiner rejects Claims 11, 13, 17, 26, 34 and 38-39 on the same rationale as Claim 1. It is the position of the Applicant that Claims 11, 13, 17, 26, 34 and 38-39 are allowable for at least the reasons discussed above with respect to Claim 1. Further:

Regarding Claim 17,

Claim 17 recites:

17. (Original) *A computer network comprising:*

a computing system configured to access a component of the computer network using a domain name;

a caching server including a first data structure configured for translating the domain name into DNS information, and means for examining the first data structure in a time that is essentially constant as a function of a number of labels comprising the domain name; and

a second data structure configured for translating the domain name into DNS information.

First, Claim 17 includes “*a first data structure,*” and a separate “*second data structure.*” It is the position of the Applicant that the cited art does not teach both “*a first data structure*” and a separate “*second data structure.*” Other than the hash table of Singh, the Examiner’s comments do not appear to point out any other data structure. Assuming for the sake of argument, that the first data structure is the “hash table” of Singh, the Applicant is unable to identify a teaching of a separate second data structure within Fletcher. Rather, Fletcher appears to be silent with regard to data structures. The Applicant, therefore, requests that the Examiner specifically point out teaching of both “*a first data structure,*” and a separate “*second data structure,*” or allow Claim 17.

Second, Claim 17 includes the limitations “*a caching server including a first data structure configured for translating the domain name into DNS information.*” The Applicant respectfully points out that in the combination of Fletcher and Singh suggested

by the Examiner, the cache of Fletcher would include the interface ID of Singh in a hash table. As illustrated in FIG. 6 of Singh, the interface ID is used as an index to a "Port." It is unclear to the Applicant how a hash table consisting of a 64-bit interface ID and a port could be "*configured for translating the domain name into DNS information.*" An interface ID and a port do not appear to be of much use in such a translation. It is, therefore, the position of the applicant that the combination of Fletcher and Singh suggested by the Examiner would not include "*a first data structure configured for translating the domain name into DNS information,*" as recited in Claim 17. The Applicant, therefore, requests that the Examiner specifically point out such a teaching within the cited art or allow Claim 17, and Claim 18, which depends therefrom.

Third, Claim 17 recites "*means for examining the first data structure in a time that is essentially constant as a function of a number of labels comprising the domain name.*" It is the position of the Applicant that the cited art does not teach these limitations. As pointed out above the combination suggested by the Examiner would include the interface ID of Singh in the cache of Fletcher. However, even assuming for the sake of argument that the interface ID represented a "label," having one label within a hash table would not result in "*a time that is essentially constant as a function of a number of labels comprising the domain name,*" as recited in Claim 17. Each label would have to be searched for separately. Therefore, the Applicant specifically requests that the Examiner explain how the combined teachings of Fletcher and Singh teach "*comprising the domain name,*" or allow Claim 17, and Claim 18, which depends therefrom.

Regarding Claim 39,

39. (Original) *The method of claim 33, wherein the answer cache is configured to store answer information and the referral cache is configured to store referral information, and the answer cache and the referral cache have different data structures.*

5

Claim 39 includes the limitation “*the answer cache and the referral cache have different data structures.*” This limitation is not found in Claim 1 and is not addressed by the Examiner. As stated in *Diamond v. Diehr*, “[w]hen evaluating the scope of a claim, every limitation in the claim must be considered,” (MPEP § 2106.II.C). It is, therefore, the position of the Applicant that the Examiner has not presented a *prima facie* case for the rejection of Claim 39. The Applicant requests that the Examiner specifically point out teachings of all of the limitations of Claim 39 in the cited art, or allow Claim 39.

10

The Applicant also believes that Claim 39 is allowable for at least the same reasons as Claim 33 from which it depends.

15

Regarding Claim 2, the Applicant believes that Claim 2 is allowable for at least the reasons discussed herein with respect to Claim 1.

Regarding Claims 25, 29 and 35,

The Applicant believes that Claim 25 is allowable for at least the reasons discussed herein with regard to Claim 1, and Claim 19 from which it depends.

20

The Applicant believes that Claim 29 is allowable for at least the reasons discussed herein with regard to Claim 1, and Claim 26 from which it depends.

The Applicant believes that Claim 35 is allowable for at least the reasons discussed herein with regard to Claim 1, and Claim 33 from which it depends.

Regarding Claim 19,

Claim 19 recites:

19. *A method of determining DNS information, the method comprising:*
receiving a request for DNS information corresponding to a domain name;
5 *examining an answer cache for answer information, the answer cache including a*
hash table configured to store the answer information or to store a pointer
to the answer information; and
searching a tree data structure if the DNS information is not found by examining
10 *the answer cache.*

10 The Examiner rejects Claim 19 on the same rationale as the rejection of Claims 1
and 2. However, it is the Applicant's position that the Examiner has not pointed out
teaching of "*examining an ... answer cache including a hash table*" in combination with
"*searching a tree data structure if the DNS information is not found by examining the*
15 *answer cache,*" as recited in Claim 19. As discussed above in reference to Claim 1, the
cited art does not appear to teach an examination of "*an answer cache*" in addition to a
separate search. Therefore, the cited art does not teach, "*examining an answer cache*" in
combination with "*searching a tree data structure.*" Likewise, as discussed above in
reference to Claim 17, the cited art does not appear to teach both a first data structure and
20 a second data structure. Therefore, the cited art does not teach both an "*answer cache*
including a hash table" and "*a tree data structure.*"

Further, it is the position of the Applicant that Fletcher does not teach an answer
cache including a "*hash table*" in combination with a "*tree data structure*" that is
searched "*if DNS information is not found by examining the answer cache.*" In paragraph
25 [0030], Fletcher does teach that a "referral" is issued if "requested data" is not found in a
query. However, assuming for the sake of argument that the query of paragraph [0030]
teaches "examining an answer cache", the Applicant is unable to find any teaching

(specific or inherent) within either Fletcher or Singh that issuing a “referral” involves “*searching a tree data structure*,” as recited in Claim 19. The Applicant, therefore, requests that the Examiner specifically point out these and the other limitations of Claim 19, or allow Claim 19, and those claims that depend therefrom.

5 Further, the Applicant believes that Claim 19 is allowable for at least the reasons discussed herein with respect to Claims 1, 11 and 22.

Regarding Claim 3,

Claim 3 recites:

10 3. *The caching server of claim 1, wherein the flat data structure includes pointers to a tree data structure.*

In rejecting Claim 3, the Examiner states:

15 Fletcher and Singh teaches the claimed invention as described above and furthermore, Fletcher teaches that when the requested address information is not found at the terminal, the query from the local host is forwarded to the communication network (i.e. to the remote hosts) which stores the requested information (e.g. see paragraph [0008]). Therefore, the pointer/link has to be inherently stored/present in the local terminal cache that points to the remote hosts for the requested information. Fletcher teaches the further limitation of pointers
20 pointing to a tree data structure (e.g. see paragraph [0005]).

The Applicant traverses this statement.

First, the Examiner appears to read the word “link” into Claim 3. Claim 3 does not recite “link.” Claim 3 recites “*the flat data structure includes pointers.*”

25 Second, the hash table of Singh, which the Examiner suggests teaches “*the flat data structure*,” is taught by Singh to include an “interface ID” used to index a “Port” and nothing more (FIG. 6). Neither an interface ID nor a port are a “*pointer.*” A person of ordinary skill in the art would not consider a pointer to be an “interface ID” or something “that points to the remote hosts” as suggested by the Examiner. For example, according

to The American Heritage® Dictionary of the English Language, Fourth Edition a “pointer” is a “variable that holds the address of a core storage location.” A remote host would not be considered a “core storage location.” Likewise, www.wikipedia.org defines “pointer” as “a programming language datatype whose value refers directly to (‘points to’) another value stored elsewhere in the computer memory using its address.” Thus, a pointer is internal to “the computer.” Neither of these exemplary definitions would be applicable to an “interface ID,” “Port” or “remote hosts” as suggested by the Examiner.

Third, it is the position of the Applicant that the cited art does not teach “pointers to a tree data structure.” The “interface ID” of Singh appears to be an identity number related to a network interface “Port” (paragraph [0028] and FIG. 6) rather than being related to “a tree data structure,” as recited in Claim 3. Specifically, as illustrated in FIG. 6, the interface ID is an index to a “Port.” Thus, the “interface ID” is not a pointer and does not point to “a tree data structure.” While the Examiner suggests that a tree data structure is taught in paragraph [0005] of Fletcher, the Applicant is unable to find any such teaching. There are no references to data structures in paragraph [0005]. Thus, the Applicant is unable to identify teachings of either “*pointers*” or “*pointers to a tree data structure*” as suggested by the Examiner. The Applicant requests that the Examiner specifically point out how the cited text teaches these limitations, or allow Claim 3.

Finally, the Examiner’s suggestion that “the pointer/link has to be inherently stored/present in the local terminal cache” is an improper inference because a pointer does not necessarily flow from the cited teachings. For example, a Universal Resource Locator (URL) rather than a “pointer” could be used in the system of Fletcher. As discussed above, an argument of inherency is only proper when the inference

“necessarily” flows from the cited teachings. Here, a pointer does not necessarily flow from the cited art because Fletcher may use, for example, a URL instead.

For at least the above reasons, the Applicant believes that Claim 3 is allowable. The Applicant further believes that Claim 3 is allowable for at least the reasons discussed
5 above with respect to Claim 1, from which it depends.

Regarding Claims 5, 20-23 and 27-28:

The Examiner rejects Claims 5, 20-23 and 27-28 “on the same rationale as the rejection of claim 3.” The Applicant believes that Claims 5, 20-23 and 27-28 are allowable for the reasons discussed above with respect to Claim 3. Furthermore:

10 **Regarding Claim 5,**

Claim 5 recites:

5. (Original) *The caching server of claim 1, wherein the flat data structure includes pointers to a tree data structure, and the tree data structure is included in the referral cache.*

15

Claim 5 further includes the limitations that “*the tree data structure is include in the referral cache.*” The Examiner does not appear to address these limitations. The Applicant therefore requests that the Examiner specifically point out teachings within the cited art of a “*flat data structure*” through which “*answer information is accessed*” and
20 that “*includes pointers to a tree data structure ... included in the referral cache*” or allow Claim 5.

The Applicant further believes that Claim 5 is allowable for at least the same reasons as Claim 1, from which it depends.

Regarding Claim 23,

Claim 23 recites:

23. (Original) *The method of claim 19, wherein the tree data structure is configured to store pointers to referral data.*

5

Claim 23, includes the additional limitations that “*the tree data structure is configured to store pointers to referral data.*” The Examiner does not appear to address these limitations. The Applicant, therefore, requests that the Examiner specifically address these limitations, or allow Claim 23.

10

The Applicant further believes that Claims 20-23 and 27-28, are allowable for at least the reasons discussed with respect to Claims 19 and 26, from which they depend.

Regarding Claim 4,

Claim 4 recites:

15

4. (Original) *The caching server of claim 1, wherein the flat data structure includes pointers to a tree data structure, and the tree data structure is configured to store answer information and referral information.*

20

In rejecting Claim 4, the Examiner states “Hudson [Fletcher] teaches that the tree data structure (i.e. the hierarchical structure) is configured to store answer information and referral information (e.g. see paragraphs [0005]-[0006]).” The Applicant traverses this statement.

25

First, while paragraph [0005] does refer to a “hierarchical structure,” this reference is a characterization of a relationship between servers and not to a hierarchical data structure, much less a “*tree data structure*” as recited in Claim 4. The text cited by the Examiner teaches, “[t]he distributed database is implemented by storing various portions of the database across multiple servers in a hierarchical structure.” This structure includes “root” servers, “authoritative” servers, and “non-authoritative” servers,

as would be understood by one of ordinary skill in the art. A “*hierarchical data structure*,” as recited in Claim 8 is, therefore, not taught in the text cited by the Examiner.

Second, a “*tree data structure*” cannot be inferred from a “hierarchical structure.”

There are hierarchical structures that are not “tree data structure[s]” (e.g., a linear
5 hierarchical structure) and thus, the limitations of Claim 4 would not “necessarily flow”
from a teaching of a hierarchical data structure, even if such a data structure were taught
in the cited art. The Applicant, therefore, requests that the Examiner specifically point
out the limitations of Claim 4 within the cited art, or allow Claim 4.

The Applicant further believes that Claim 4 is allowable for at least the same
10 reasons as Claim 1, from which it depends.

Regarding Claim 6,

Claim 6 recites:

6. *The caching server of claim 1, wherein the caching server is also an authoritative
server.*
15

With regard to Claim 6, the Examiner states “Hudson [Fletcher] teaches that the
caching server (i.e. the name server, 101 Fig. 1) is also an authoritative server, i.e. a
server which has the desired information (e.g. see Fig. 1).” The Applicant traverses the
Examiner’s definition of “authoritative server.” It is the position of the Applicant that a
20 person of ordinary skill in the art would not define “authoritative server” as merely “a
server which has the desired information,” as suggested by the Examiner. To the
contrary, a non-authoritative server may also have such information. For example,
further characterizations of the term “authoritative server” include:

25 There are two common types of DNS servers: the Authoritative DNS Name
Server and The Non-Authoritative DNS Name Caching Server. ... An
Authoritative DNS server is the authoritative source for all DNS requests made

for a designated zone or domain.
(<http://www.more.net/technical/dns/overview.html>)

Thus, the Applicant is unable to find any teaching within the cited art that the server of
5 Fletcher is an authoritative server rather than a non-authoritative server. The Applicant,
therefore, requests that the Examiner provide support for the suggestion that a teaching of
“a server which has the desired information” teaches an “*authoritative server*,” as recited
in Claim 6, or allow Claim 6.

The Applicant further believes that Claim 6 is allowable for at least the same
10 reasons as Claim 1, from which it depends.

Regarding Claims 7,

Claim 7 recites:

7. (Original) *The caching server of claim 1, wherein the caching server is also a web
15 server.*

In rejecting Claim 7, the Examiner states “Hudson [Fletcher] teaches that the
caching server (i.e. the name server, 101 in Fig 1) is also a web server, i.e. the DNS
server (e.g. see Fig. 1).” The Applicant traverses this statement. In paragraph [0027]
Fletcher characterizes element 101 of FIG. 1 as “a name server” configured to handle
20 resolution traffic using UDP or TCP. However, the Examiner is unable to find any
teaching that this “name server” is a “*web server*” as recited in Claim 7. A name server
may be, for example, part of a local network. Thus, a teaching of a “name server” does
not necessarily teach a “web server.” The Applicant requests that the Examiner
specifically point out teaching that the “name server” of Fletcher is a “*web server*,” or
25 allow Claim 7.

The Applicant further believes that Claim 7 is allowable for at least the same reasons as Claim 1, from which it depends.

Regarding Claim 8,

Claim 8 recites:

- 5 8. (Original) *The caching server of claim 1, wherein the referral cache is further configured to store the referral information in a hierarchical data structure.*

10 In rejecting Claim 8, the Examiner states “Hudson [Fletcher] teaches that the referral cache is further configured to store the referral information in a hierarchical data structure (e.g. see paragraphs [0005]-[0006]).” The Applicant traverses this statement. The text cited by the Examiner teaches, “[t]he distributed database is implemented by storing various portions of the database across multiple servers in a hierarchical structure.” This text concerns a hierarchical relationship between servers and not a hierarchical data structure. This structure includes “root” servers, “authoritative” servers, 15 and “non-authoritative” servers, as would be understood by one of ordinary skill in the art. A “*hierarchical data structure*,” as recited in Claim 8, is therefore, not taught within the text cited by the Examiner. The Applicant, therefore, requests that the Examiner point out teachings of all the limitations of Claim 8, or allow Claim 8.

20 The Applicant further believes that Claim 8 is allowable for at least the same reasons as Claim 1, from which it depends.

Regarding Claims 31 and 41, it is the Applicant’s position that Claims 31 and 41 are allowable for at least the reasons discussed herein with respect to Claims 26 and 40, from which they depend, respectively. Further, Claims 31 and 41 are allowable for at least the reasons discussed herein with respect to Claim 8.

Regarding Claims 9, 12, 14, 15, 16, 18, 24, 30, 32 and 42,

The Applicant believes that Claims 9, 12, 14, 15, 16, 18, 24, 30, 32 and 42 are allowable for at least the reasons discuss herein with respect to the claims from which they depend, respectively.

5 Regarding Claim 33,

Claim 33 recites:

33. *A method of storing data in a cache, the method comprising:*
requesting DNS information;
receiving data in response to the request;
10 *classifying the response received; and*
storing the data received in either a referral cache or an answer cache based on
the classification.

In reference to Claim 33, the Examiner states “Fletcher teaches ... classifying the
15 response received, and storing the data received in either a referral cache or an answer
cache based on the classification (e.g. see paragraphs [0005] and [0008]).” The
Applicant traverses this statement.

First, it is unclear to the Applicant which parts of Fletcher are thought by the
Examiner to teach, “*classifying the response received.*” Neither paragraphs [0005] nor
20 [0008] refer to classification of a response. The Applicant is, thus, unable to identify any
teaching within Fletcher of the teachings suggested by the Examiner. The Applicant,
therefore, requests that the Examiner specifically point out teaching of “*classifying the*
response received,” or allow Claim 33, and those claims that depend therefrom.

Second, it is unclear to the Applicant which parts of Fletcher are thought by the
25 Examiner to teach, “*storing the data received in either a referral cache or an answer*
cache based on the classification.” As discussed in regard to Claim 1, Fletcher does not
explicitly teach both an answer cache and a referral cache, and such a teaching cannot

properly be inferred. As discussed in the paragraph above, the Applicant is also unable to identify any teaching of a “classification” within the cited art. Thus, the Applicant is unable to identify numerous limitations of Claim 33 in the cited text. The Applicant, therefore, requests that the Examiner specifically point out teaching of “*storing the data*
5 *received in either a referral cache or an answer cache based on the classification,*” or allow Claim 33, and those claims that depend therefrom.

Regarding Claims 36 and 37, the Applicant believes that Claims 36 and 37 are allowable for at least the same reasons as Claim 33, from which they depend.

Regarding Claim 40,

10 **Claim 40 recites:**

40. *A method of caching DNS information, the method comprising:*
requesting DNS information;
receiving data in response to requesting DNS information;
classifying the response received as an answer response or a referral response;
15 *storing the response received in either a referral cache or an answer cache based*
on the classification, the answer cache including a flat data structure;
receiving a request for DNS information corresponding to a domain name;
examining the answer cache to find answer information, responsive to the
20 *received request; and*
responsive to the examination of the answer cache, searching the referral cache.

Regarding Claim 40, the Examiner states “Claim 40 is also rejected based on the same rationale as the rejection of claims 1 and 33.” The Applicant believes that Claim 40 is allowable for at least the same reasons as Claims 1 and 33.

25 In addition, Claim 40 includes limitations not found in either Claims 1 or 33. For example, neither of these claims includes “*classifying the response received as an answer response or a referral response.*” The Examiner has not addressed these limitations and, thus, not made a *prima facie* case for the rejection of Claim 40 under §103. Further, the Applicant is unable to find any teaching of these limitations within Fletcher. The

Applicant, therefore, respectfully requests that the Examiner specifically point out these limitations within the cited art, or allow Claim 40, and those claims that depend therefrom.

Rejections under 35 USC § 103

- 5 **Claim 10 is rejected under 35 U.S.C. 103(a) as being unpatentable over Fletcher in view of Singh, further in view of Ramanathan et al. (USPN: 6,182,136).**

Regarding Claim 10,

Claim 10 recites:

10. The caching server of claim 1, wherein the DNS information includes an MX record.

- 10 Regarding Claim 10 the Examiner admits that Fletcher and Singh do “not clearly teach that the DNS information includes a Mx Record,” and cites Ramanathan et al. as teaching the inclusion of an Mx record in DNS information. The Applicant traverses this statement.

- 15 First, as discussed above in regard to Claim 1, Fletcher and Singh do not teach all the limitations of Claim 1, from which Claim 10 depends, as suggested by the Examiner. Thus, even in combination with Ramanathan et al., not all of the limitations of Claim 10 are taught by the cited art.

- 20 Second, in the combination of Fletcher and Singh, the Examiner appears to suggest that examining a mere “interface ID” (Singh) can be used to translate a domain name in to DNS information. In Claim 10, this DNS information is recited as including “an MX record.” It is not clear to the Applicant how an examination of a mere “interface ID” would be sufficient to translate a domain name into an Mx record. This combination does not appear to have a reasonable likelihood of success. Thus, the proposed

combination does not appear to satisfy the second requirement for a *prima facie* case of obviousness as discussed in In re Vaeck, *supra*. The Applicant, therefore, requests that the Examiner provide a detailed explanation of how such a combination would work, or allow Claim 10.

5 Third, as a motivation to combine the cited art, the Examiner appears to merely state an advantage of the suggested combination. As discussed with respect to Claim 1, the requirements of a *prima facie* case under §103 include “some suggestion or motivation” for the combination, not that the combination would produce a benefit. An invention can be non-obvious while still producing a benefit. The advantage must be
10 apparent without the benefit of hindsight. The Applicant, therefore, requests that the Examiner provide “some suggestion or motivation, either in the references themselves or in the knowledge generally available to one of ordinary skill in the art, to modify the reference or to combine reference teachings” as required to establish a *prima facie* case under §103, or allow Claim 10.


15 The Applicant further believes that Claim 10 is allowable for at least the reasons discussed herein with respect to Claim 1, from which it depends.

Applicant believes that all pending claims are allowable and respectfully requests that the Examiner issue a Notice of Allowance. Should the Examiner have questions, the Applicant's undersigned representative may be reached at the number provided.

5

Respectfully submitted,
Andreas Gustafsson

10 Date: May 12, 2006



Steven M. Colby, Ph.D. Reg. No. 50,250
Carr & Ferrell *LLP*
2200 Geng Rd.
Palo Alto, CA 94303
Phone (650) 812-3424
Fax (650) 812-3444

15